arXiv:1610.04297v1 [stat.AP] 14 Oct 2016

On distribution free testing in Bernoulli trials with distributions depending on covariates

Leigh A Roberts*
Department of Economics and Finance, Victoria University of Wellington

October 17, 2016

Abstract

In a sequence of independent Bernoulli trials, generated by a common distribution but with the outcome of each trial depending on a covariate varying over the sample, we carry out distribution free goodness of fit tests by considering the sample as a sequence of subgroups of trials. Within each group, the empirical process is 'rotated' by means of unitary transformations to a common stochastic process defined on a given reference distribution. For ungrouped trials, or groups of size one, the transforms are valid when the null hypothesis is completely specified. In the more usual case in which the null hypothesis is specified only up to unknown parameters, however, it is necessary to group the data and apply a unitary transform to each subgroup. For group sizes of 2, 3 and 4, we apply the Kolmogorov Smirnov statistic to investigate power to reject incorrect null hypotheses.

Keywords: grouping binomial trials, logistic distribution, unitary transform, Kolmogorov-Smirnov statistic

*Thanks to Estate Khmaladze for encouragement to study this area of statistics and for comments on early drafts of this paper. Responsibility for errors naturally remains with the author.

1 Introduction

In seeking to transform statistical goodness of fit tests into a distribution free format, consider first a centred empirical process having the basic form:

$$\frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left[\mathbb{1}_{\{\xi_i \leq y\}} - F_\theta(y) \right] \tag{1}$$

in which $\mathbb{1}$ is the indicator function; $\{\xi_i\}_{i=1}^n$ is the random sample of independent and identically distributed drawings; and $F_\theta(.)$ the distribution function in question. For fixed parameter θ, the process in (1) is asymptotically a Brownian bridge in time F_θ. Standard goodness of fit tests such as the Kolmogorov Smirnov statistic may then be applied independently of the nature of the assumed distribution F_θ, giving rise to so-called omnibus goodness of fit tests of the hypothesised or null distribution.

The extension of this testing procedure to the case in which the parameter θ needs to be estimated from the data has been resolved by Khmaladze in two distinct ways. As long ago as 1981, he proposed an alternative centring to the empirical process to produce a process with independent increments, suitable for testing goodness of fit through a limiting process of a Brownian motion rather than a Brownian bridge: see Khmaladze (1981); and Koul and Swordson (2011), together with references therein. Recently Khmaladze has developed a method for transforming a process defined on one space to a process defined on another space; choosing to transform a given empirical process to an unrelated reference process enables one to test goodness of fit using a common distribution, giving rise to a generic or omnibus test. By transforming the data in this way, one can test whether data is generated by the hypothesised distribution by testing for whether the transformed process is likely to have been generated by the common reference distribution, chosen to be relatively simple.

The general theory for this latter case appears in Khmaladze (2016*b*), with Khmaladze (2013) specialising to the more transparent case in which the distribution F is discrete.

More recently Khmaladze (2016a) has extended the methodology to the case in which the sample contains independent drawings from a common distributional family, but the distributions for individual sample trials depend on a covariate, possibly different for each sample point: the sample values are independent, but no longer identically distributed.

We apply the theory as set out in Khmaladze (2016a) to a sample of independent Bernoulli trials, considered alternately as a sequence of subgroups of sizes 1, 2, 3 and 4, labelled as singles, doubles, triples and quadruples respectively. Empirical processes within each subgroup, including groups of size one, are transformed to a common reference process. The unitary transformations utilised in this paper can be applied to individual Bernoulli trials when the parameter θ is fixed, but the data must be grouped to allow distribution free goodness of fit tests in the more usual case when the parameter is to be estimated from a given family of distributions.

While the application to Bernoulli trials may seem relatively simple, the situation is one commonly occurring in practice, and our results here illustrate well the utility of Khmaladze's transformation of one empirical process to a process following a generic unrelated distribution: computational requirements are modest, and the power of the Kolmogorov Smirnov statistic to reject the null when it is false is high.

The segmentation of Bernoulli trials into groups may also seem artificial, but gathering trials into clusters is frequently a sensible approach to take in modelling: a simple example could be to group the results for an individual in a recent survey with those obtained from the same individual a year earlier. Should the hypothesis of independence of trials be rejected, one could consider dependence within subgroups, which may involve assuming a common covariate within a group. Although in this paper we assume the single covariate to be chosen independently from one Bernoulli trial to the next, extensions to several covariates of more general types is easily encompassed within the framework utilised in

this paper.

Economists have frequently used the logistic distribution to model dichotomous choices or outcomes: see for instance Chamberlain (1984, p. 1274), Duffie and Singleton (2003, p. 77) and Maddala (1989, p. 272). More generally, binomial/logistic models also figure prominently in statistical analysis of categorical data, for which see Agresti and Kateri (2011) and McCullagh and Nelder (1989), amongst many others. This provides the motivation for our choice of this distribution for the illustration of the methodology given below.

In §2 we set out the form of the empirical process underpinning our work in this paper. Khmaladze's seminal transformation of a stochastic process from one distribution to another is discussed, and applied to subgroups of sizes 1, 2, 3 and 4. We shall see that, although these transforms become more complicated as we allow for more parameters to be estimated from the one distributional family, computation is straightforward whatever the size of the group because calculations can be effected via a sequence of self-adjoint unitary transformations of a simple given form. Recently Nguyen (2016) has demonstrated that parameters of dimension 10 or more can be easily handled.

Simulation results are presented in §3. Three distributions are used for generating data, firstly to verify that they all map to the same common target distribution when the correct null hypothesis is assumed, so that we do in fact have a distribution free test; and secondly to use the Kolmogorov Smirnov statistic to test the power of our methodology to reject false null hypotheses, when the null hypothesis assumes the data to be generated by the logistic distribution. A short conclusion summarises and lists ideas for future work.

2 Preamble

2.1 The empirical process

Consider a sample of nm independent Bernoulli trials, subdivided into n groups each containing m trials. The jth trial within the ith subgroup is associated with a covariate X_{ij}, and produces a result denoted ξ_{ij}, for $1 \leq i \leq n$, and $1 \leq j \leq m$. The nm covariates X_{ij} are mutually independent; and each variable ξ_{ij} assumes the value of 0 or 1, the former labelled a 'failure' and the latter a 'success'.

We do not model the behaviour of the covariate x, and work with a joint distribution of the form

$$P_\theta(dx, dy) = P_{x,\theta}(dy)H(dx) \tag{2}$$

in which the distribution H is unspecified. For the probability measure P_θ, then, set $p_{X_{ij},\theta}(y_{ij}) = p(\xi_{ij} = y_{ij}\big|X_{ij}, \theta)$ for $y_{ij} = 0, 1$; and let $x_i = \{x_{ij}\}_{j=1}^m, y_i = \{y_{ij}\}_{j=1}^m, X_i = \{X_{ij}\}_{j=1}^m$ and $\xi_i = \{\xi_{ij}\}_{j=1}^m$. Denoting by A and B Borel sets in the respective spaces, we define the following centred empirical process:

$$\alpha_n^P(A, B, \theta) = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left[\mathbb{1}_{\{\xi_i \in B\}} - P_\theta(B\big|X_i) \right] \mathbb{1}_{\{X_i \in A\}} \tag{3}$$

Later we shall see that it is more convenient to parametrise processes by functions rather than sets: see Khmaladze (2016*b*, p. 1), *i.a.*, for a brief discussion on the alternative ways in which to parametrise processes. Considering functions $\phi(x, y)$ in $L_2(P_\theta)$, the space of square integrable functions with respect to the measure P_θ, we rewrite (3) in the form

$$\alpha_n^P(\phi, \theta) = \int\int \phi(x, y)\alpha_n^P(dx, dy, \theta) \tag{4}$$

2.2 Transforms to a different space

We wish to transform processes defined on $L_2(P_\theta)$ to processes defined on say $L_2(Q_\theta)$, where by analogy to (2) we have

$$Q_\theta(dx, dy) = Q_{x,\theta}(dy)H(dx)$$

Consider functions $\phi \in L_2(P_\theta)$ and $\psi \in L_2(Q_\theta)$. Provided that P_θ and Q_θ are equivalent, define

$$\ell_{x,\theta}(dy) = \sqrt{\frac{dQ_{x,\theta}(dy)}{dP_{x,\theta}(dy)}} \tag{5}$$

so that $\psi \in L_2(Q_\theta)$ if and only if $\ell\psi \in L_2(P_\theta)$. The functions ϕ and ψ are square integrable in y with respect to $P_{x,\theta}$ and $Q_{x,\theta}$ respectively, for $H-$almost all x, and are thus square integrable with respect to H. The precise form of H being unspecified, it is convenient to assume that the functions ϕ and ψ are bounded in x; since y assumes but finitely many values, the functions are bounded in both x and y.

Our intention is to simplify the conditional $Q_{x,\theta}(y)$ measure so that the numerator of the quantity under the square root in (5) depends on neither x nor θ: we may then write

$$\ell_{x,\theta}(y) = \sqrt{\frac{\prod_{j=1}^m q^{1-y_j}(1-q)^{y_j}}{\prod_{j=1}^m p_{x,\theta}^{1-y_j}(1-p_{x,\theta})^{y_j}}} \tag{6}$$

with q and $p_{x,\theta}$ denoting probabilities of failure, and $y = 1$ indicating success.

2.3 The unitary transforms

For functions $\eta, \zeta \in L_2(P_\theta)$ of unit norm, consider the unitary operator acting on $L_2(P_\theta)$:

$$\begin{aligned} U_{\eta,\zeta}(\phi) &= \phi - \frac{2}{||\eta-\zeta||^2}(\eta-\zeta) < \eta-\zeta, \phi > \\ &= \phi - \frac{1}{1- < \eta, \zeta >}(\eta-\zeta) < \eta-\zeta, \phi > \end{aligned} \tag{7}$$

in which the norm and the inner product are defined with respect to the P_θ measure:

$$< \phi_1, \phi_2 > = \sum_y \phi_1(x,y)\phi_2(x,y)p_{x,\theta}(y)$$

The operator $U_{\eta,\zeta}$ interchanges η and ζ; more precisely it 'rotates' vectors in the 2 dimensional subspace spanned by η and ζ, and leaves vectors orthogonal to that subspace unchanged.

The transformation of the empirical process on $L_2(P_\theta)$ to that on $L_2(Q_\theta)$ is effected as follows. From Khmaladze (2016*a*) and (4),

$$\alpha_n^Q(\psi,\theta) = \alpha_n^P(U\ell\psi,\theta) \tag{8}$$

in which U is a unitary matrix assuming a specific form. In the simplest case for which θ is assumed known, $U = U_{\mathbf{1},\ell}$ in (7), with the function $\mathbf{1}$ defined to be identically unity, and $\ell_{x,\theta}$ abbreviated to ℓ.

2.4 The Kolmogorov Smirnov goodness of fit test

We wish to carry out a goodness of fit test as to whether or not the data could reasonably have been generated by P_θ, for some value of θ, assumed for the moment to be known. A standard means of doing this would be to employ the Kolmogorov Smirnov (KS) statistic, defined from (4) as

$$\max_{\phi\in\Phi} |\alpha_n^P(\phi,\theta)| \tag{9}$$

where the maximum is taken over all functions ϕ in a class Φ of functions in $L_2(P_\theta)$. The point of transforming to the Q space, however, is that we may instead employ the KS statistic in the Q space:

$$\max_{\psi\in\Psi} |\alpha_n^Q(\psi,\theta)| \tag{10}$$

A suitable class Ψ is generated by indicator functions in both x and y: the former are conventional, since x is a continuous variable, but the latter slightly less so.

As an illustration, consider the doubles $m = 2$, the outcome of which for the ith sample point is one of $00, 01, 10$ and 11, regarded as binary numbers in the order given. Centring the second indicator function results in the following family of functions Ψ, with $x_0 \in \mathbb{R}$; and y_0 is one of $(00, 01, 10, 11)$, with the corresponding $y_0^* = (1, 2, 3, 4)$:

$$\psi_{x_0,y_0}(x, y) = \mathbb{1}_{\{x \le x_0\}} \left[\mathbb{1}_{\{y \le y_0\}} - \frac{y_0^*}{4} \right] \tag{11}$$

The extension to triples and quadruples is immediate.

The covariates used for the experiments reported in this paper were drawn from the uniform distribution between 0 and 2, and were chosen independently for each Bernoulli trial. For $m = 2$ and 3, the grid points x_0 were chosen to subdivide the interval (0,2) into 100 subintervals: for $m = 2$, the 2 dimensional x_0 values gave a grid with 10,000 cells, and for triples the grids had 1,000,000 cells; but the number of subintervals along each axis was reduced to 20 for $m = 4$ in order to reduce the requisite computing time.

2.5 The nature of the operator U for unknown θ

When the data is generated in P_θ, and the parameter θ is fixed and known, the transformation from α_n^P to α_n^Q is given in (8), with $U = U_{\mathbf{1},\ell}$.

When θ is a single parameter to be estimated from the data, U is again a unitary transform in $L_2(P_\theta)$, but of a more complicated kind. Define (normalised) score functions $a_1 = \Gamma_{xP}^{-1/2} \dot{p}/p$ and $b_1 = \Gamma_{xQ}^{-1/2} \dot{q}/q$, with Γ_{xP} defined in (17) and (13), and Γ_{xQ} defined analogously. We still require U to map ℓ to $\mathbf{1}$, as previously; but now in addition we require U to map ℓb_1 to a_1. This can be arranged via a sequence of operators of the form (7).

For let $\widetilde{\ell b_1} = U_{\mathbf{1},\ell}(\ell b_1)$; then the required U is given as

$$U = U_{a_1,\widetilde{\ell b_1}}(U_{\mathbf{1},\ell}) \tag{12}$$

with θ replaced by its estimated value (Khmaladze, 2016*b*). The score functions are discussed further in the next section.

When the unknown parameter θ is of higher dimension, say K, there will be further vectors a_i, b_i as components of the score function, and U will need to map ℓb_i to a_i for $i = 1, 2, \ldots, K$. The required U may be constructed as a succession of operators of the type (7), just as above in (12) when θ contains but a single component.

The particular case of $m = 1$ deserves attention. The 'rotations' are now in a $2^m = 2$ dimensional space, and the ability to impose a second constraint on the transform U in (8) cannot be assumed. In fact, when $m = 2$ it is easily seen that $U_{a_1,\widetilde{\ell b_1}}$ interchanges a_1 and $-a_1$, while leaving the unit vector $\mathbf{1}$ unchanged.

An immediate consequence is that, while unitary transforms within 'groups' of singles are adequate for goodness of fit tests for known θ, for unknown θ we must group the Bernoulli trials in order to obtain suitable transforms.

2.6 The score functions

The main task remaining before we can carry out the simulations is to find the normalised score functions a_1 and b_1, conditional upon the x values.

For a single Bernoulli trial, with covariate x, the Fisher information is

$$\Gamma_x^{(1)}(\theta) = \sum_{y=0}^{1} \left(\frac{\dot{p}_{x,\theta}(y)}{p_{x,\theta}(y)} \right)^2 p_{x,\theta}(y) \tag{13}$$

and the normalised score function a_1 is given by

$$a_1^T = \frac{1}{\sqrt{\Gamma_x^{(1)}(\theta)}} \left(\frac{\dot{p}_{x,\theta}(0)}{p_{x,\theta}(0)}, \frac{\dot{p}_{x,\theta}(1)}{p_{x,\theta}(1)} \right) \tag{14}$$

In the particular case of the logistic distribution (18)

$$\Gamma_x^{(1)}(\theta) = x^2\, p_{x,\theta}(0)\, p_{x,\theta}(1) \qquad \text{and} \qquad a_1^T = \left(-\sqrt{\frac{p_{x,\theta}(1)}{p_{x,\theta}(0)}}, \sqrt{\frac{p_{x,\theta}(0)}{p_{x,\theta}(1)}} \right) \tag{15}$$

It is natural to choose $q_{x,\theta}(0) = q_{x,\theta}(1) = 1/2$; and, with the logistic distribution as the default for the null hypotheses in the latter part of the paper, it further seems sensible in the light of (15) to specify $b_1 = (-1, 1)$. The arbitrary reference distribution Q we have chosen is a limiting case of the logistic distribution.

For doubles, with $x = (x_1, x_2)$, the score function a_1^T becomes

$$\frac{1}{\sqrt{\Gamma_x^{(2)}(\theta)}} \left(\frac{\dot{p}_{x,\theta}(00)}{p_{x,\theta}(00)}, \frac{\dot{p}_{x,\theta}(01)}{p_{x,\theta}(01)}, \frac{\dot{p}_{x,\theta}(10)}{p_{x,\theta}(10)}, \frac{\dot{p}_{x,\theta}(11)}{p_{x,\theta}(11)} \right) \tag{16}$$

with obvious extensions for groups of larger size; and in which

$$\Gamma_x^{(m)}(\theta) = \sum_{j=1}^{m} \Gamma_{x_j}^{(1)}(\theta) \qquad \text{for } m > 1 \tag{17}$$

The Q measure is taken throughout to be the uniform distribution over y, and analogous arguments to those above for $m = 1$ lead successively to the following score functions b_1^T:

$$\sqrt{2}(-1, 0, 0, 1) \qquad \frac{1}{\sqrt{3}}(-3, -1, -1, 1, -1, 1, 1, 3)$$

and

$$\frac{1}{\sqrt{4}}(-4, -2, -2, 0; -2, 0, 0, 2; -2, 0, 0, 2; 0, 2, 2, 4)$$

3 Simulation of Bernoulli trials

Our aim in running simulations is twofold. On the one hand we wish to verify that fitting a model specified in a correct null hypothesis, and transforming to the reference distribution via (8), produces the same reference distribution Q whatever the underlying data generating process. In this respect our approach mirrors that in Nguyen (2016), save that she uses a larger number of parameters and considers the omega-squared statistic as well as the KS statistic.

Our test of consistency in this sense is a visual one. For each of three models assumed to generate data, we transform from the underlying distribution to the distribution Q, and calculate the KS statistic by means of (10). Drawing new covariate values for each replication, we build up an empirical distribution function of the KS statistic. The resulting graphs are labelled 'Consistency'; and the expectation is that the resulting distribution functions will be close to each other.

The second purpose is to check on the power of the transform in (8) to pick up incorrectly specified null hypotheses. The same three distributions are used to generate data; but this time the null hypothesis is a common one, viz. that the data has been generated by the logistic distribution, the first distribution discussed below. Fitting the logistic distribution to all three data sets, we generate empirical distribution functions for the KS statistics in the same manner as previously. Labels for these graphs are 'Power', and the expectation now is that, with the incorrect null, the second and third distribution functions will be well removed from that obtained by use of the logistic distribution.

The three distributions are described first. We then look at the consistency of the mapping from the P_θ distributions to the common Q distribution, after which we examine the power to reject an incorrect null hypothesis. Sample size throughout is 96 Bernoulli

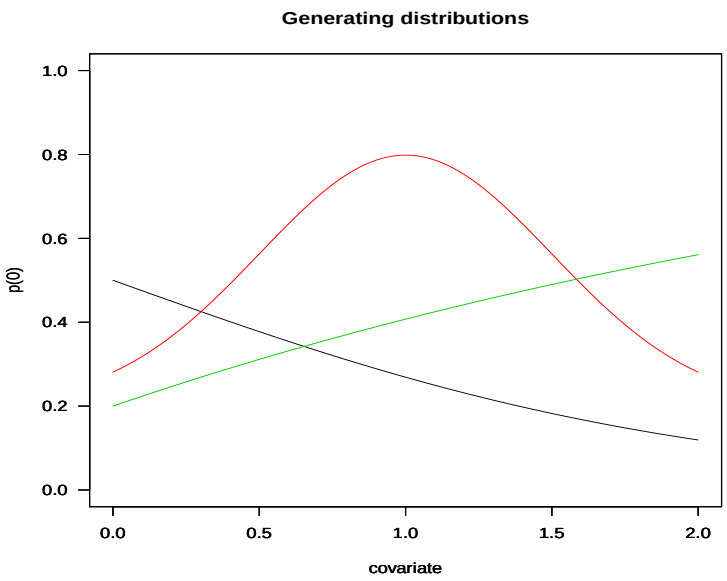

Figure 1: rg158. Data generating processes for Bernoulli trials

trials.

3.1 The three models

Covariates are chosen uniformly between 0 and 2, say x_k for $1 \leq k \leq nm$. When θ is unknown, it is estimated by maximum likelihood over the nm data points; and θ_0 is the correct value.

Figure 1 indicates several functions for mapping the covariate x to the probability of failure $p_{x,\theta}(0)$. The downward sloping solid line shows the logistic distribution, with failure probability:

$$p_{x,\theta}(0) = \frac{1}{1 + e^{\theta\, x}} \qquad \text{with } \theta_0 = 1 \tag{18}$$

which provides the first of our data generating models for the Bernoulli trials, and the

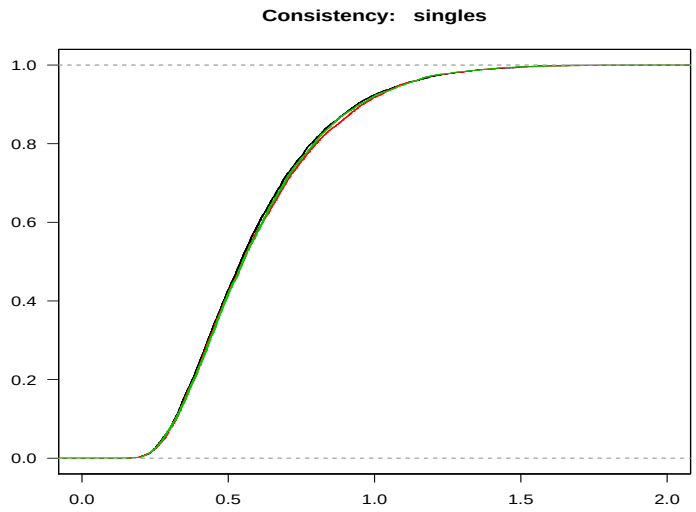

Figure 2: Consistency, known θ, 96 trials, singles, 5000 replications

default fitted model; the true parameter value $\theta_0 = 1$ has been used in Figure 1. The second data generating model is furnished by a modified normal density:

$$p_{x,\theta}(0) = \frac{3/2}{\sqrt{2\,\pi}} \exp\left[-2(x-\theta)^2\right] + 0.2 \qquad \text{with } \theta_0 = 1 \tag{19}$$

For the final data generating process we rescale the distribution function of the exponential distribution:

$$p_{x,\theta}(0) = 0.2 + 0.8(1 - e^{-\theta\, x}) \qquad \text{with } \theta_0 = 0.3 \tag{20}$$

which provides the upward sloping graph in Figure 1.

3.2 Consistency

Figures 2, 3 and 4 all give a reasonable indication of consistency when we transform the correct underlying P_θ distributions to a reference Q distribution. The first graph utilises

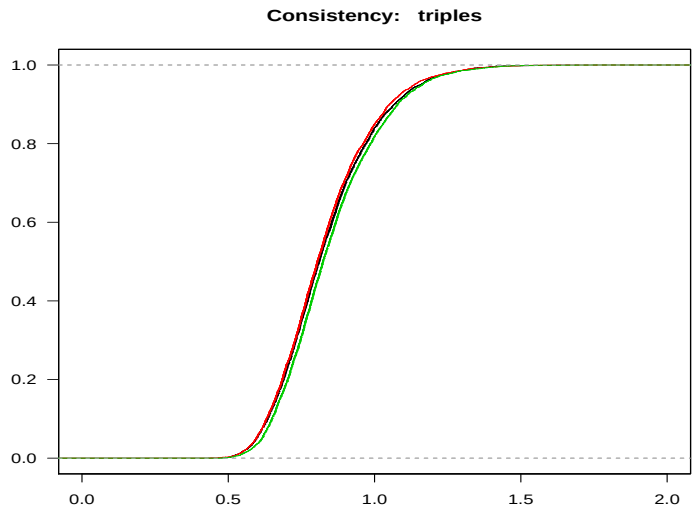

Figure 3: Consistency, unknown θ, 96 trials, triples, 5000 replications

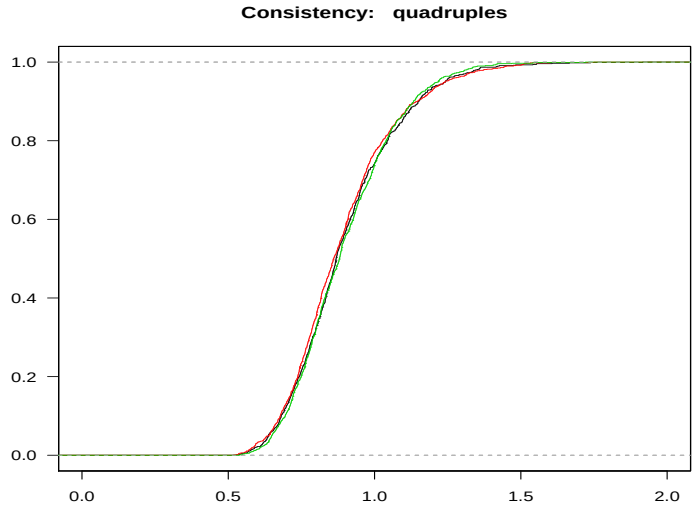

Figure 4: Consistency, unknown θ, 96 trials, quadruples, 1000 replications

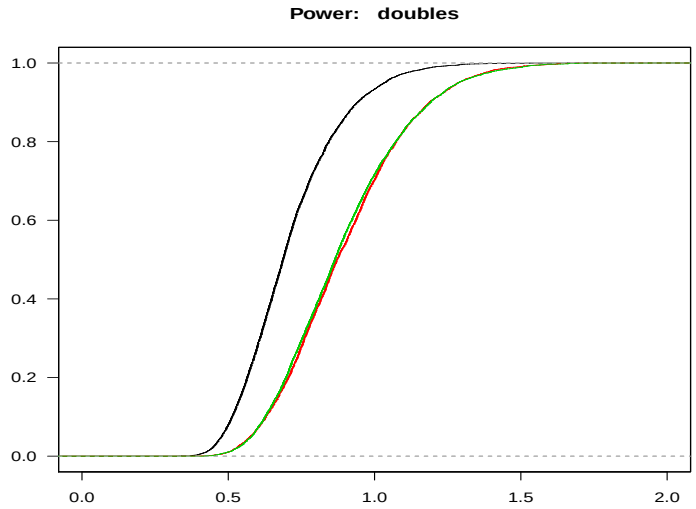

Figure 5: Power, unknown θ, 96 trials, doubles, 5000 replications

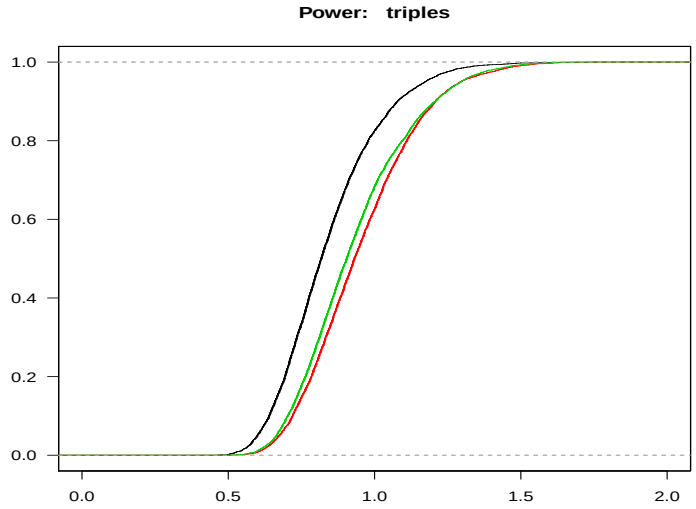

Figure 6: Power, unknown θ, 96 trials, triples, 5000 replications

the transform (8) with $U = U_{\mathbf{1},\ell}$, while the latter two apply the transform shown in (12). The first two graphs use a grid size of 100 intervals between 0 and 2, and 5000 replications; the final graph for quadruples utilises only 20 intervals and 1000 replications to economise on computing time, but is seen still to give satisfactory results.

3.3 Power

Fitting the logistic distribution (18) by maximum likelihood, Figure 5 exhibits reasonable power to reject the incorrect null hypotheses when data is grouped into pairs. The evidence for triples is not so striking, to judge from Figure 6; and, although not shown, the power for quadruples is similar to that for triples.

4 Conclusion

Despite the higher power obtained by using doubles rather than triples and quadruples for the sample sizes and data generating processes considered here, applying the unitary transforms above in higher dimensional spaces might turn out to be advantageous or even necessary under more general conditions. Dependence between Bernoulli trials may need to be modelled within large groups, by suitable alterations to the empirical process (4); and it may be the case that when several parameters need to be estimated, one would need higher dimensional spaces for the transform of stochastic processes to be effective.

While much work remains to investigate the extent to which it is desirable to work with groups of different sizes, and to apply the methodology to more general situations than Bernoulli trials, the transform of stochastic processes as set out in Khmaladze (2016*a*) is seen to be an effective tool for providing distribution free goodness of fit tests. Utilisation of the procedure is straightforward, and the computational burden modest.

References

Agresti, A. and Kateri, M. (2011), Categorical Data Analysis, *in* M. Lovric, ed., 'International Encyclopedia of Statistical Sciences', Vol. 2, Springer, pp. 206–208.

Chamberlain, G. (1984), *Panel Data*, Vol. 2 of *Handbook of Econometrics*, Elsevier, chapter 22, pp. 1247–1318.

Duffie, D. and Singleton, K. J. (2003), *Credit risk: pricing, measurement and management*, Princeton University Press.

Khmaladze, E. V. (1981), 'Martingale approach in the theory of goodness-of-fit tests', *Theory of Probability and its Applications* **26**(2), 240–257.

Khmaladze, E. V. (2013), 'Note on distribution free testing for discrete distributions', *Annals of Statistics* **41**(6), 2979–2993.

Khmaladze, E. V. (2016*a*), 'Distribution free testing for discrete distributions with covariates', *Forthcoming* .

Khmaladze, E. V. (2016*b*), 'Unitary transformations, empirical processes and distribution free testing', *Bernoulli* **22**(1), 563–588.

Koul, H. L. and Swordson, E. (2011), Khmaladze transformation, *in* 'International Encyclopedia of Statistical Science', Springer, pp. 715–718.

Maddala, G. S. (1989), *Introduction to Econometrics*, Macmillan.

McCullagh, P. and Nelder, J. A. (1989), *Generalized linear models*, 2nd edn, Chapman and Hall.

Nguyen, T. T. M. (2016), 'A new approach to distribution free tests in contingency tables', *Metrika* . doi: 10.1007/s00184-016-0596-6.